

Saurabh Kadekodi

Research Scientist (Google) — <https://cs.cmu.edu/~saukad> — saukad@google.com

RESEARCH VISION **Designing and building exascale distributed systems by replacing design decisions that are disrupted by scale with new solutions that are efficient, scalable, nimble, and robust.**

EDUCATION **Carnegie Mellon University** 2014 - 2020

Ph.D. in Computer Science

Advisors: Prof. Gregory Ganger, Prof. Rashmi Vinayak and Prof. Garth Gibson

Thesis: **Disk-Adaptive REdundancy (DARE)**: tailoring data redundancy to disk-reliability heterogeneity in cluster storage systems (*dissertation*)

Northwestern University 2012 - 2013

Masters in Computer Science

Advisor: Prof. Peter Dinda

Pune Institute of Computer Technology (University of Pune) 2005 - 2009

Bachelors in Computer Engineering

RESEARCH INTERESTS **Exascale Distributed Systems, Fault Tolerance, Data Layout and Management, File and Storage Systems, Benchmarking**

SELECT RESEARCH EXPERIENCE **Redesigning reliability in exascale storage clusters**

Conceptualized Disk-Adaptive REdundancy (DARE), in which data redundancy is dynamically tailored to observed disk failure rate heterogeneity. Analyzed over 5.3 million disks spanning 60+ makes/models from production environments of Google, NetApp, and Backblaze. Designed three DARE systems: *HeART*, *Pacemaker* and *Tiger* which provide up to 20% space-savings (worth multi-million dollars) in large-scale storage clusters while never compromising data safety.

Optimizing data management for emerging hardware and cloud

Built several data management systems including *SMRfs* – a new file system for efficient data management on Shingled Magnetic Recording (SMR) disks, *Caveat Scriptor* – a new data layout mechanism for minimizing wasted space on host-managed SMR disks, *WineFS* – a huge-page-friendly, NUMA-aware, age-resistant file system for Persistent Memory (PM). Built a *packing and indexing* layer to improve small-file ingress throughput in Amazon S3 by 60× and reduce cost by 25000×.

Building real-world benchmarking techniques

Built tools and frameworks that can synthesize real-world workloads or snapshots by examining deployed systems to promote effective and realistic evaluation. *Geriatric* accurately ages local file systems by recreating file and free space fragmentation seen after decades of usage in less than a day. Quilt (in submission) precisely reconstructs a full-resolution disk IO trace with over 95% accuracy by meticulously stitching together downsampled IO traces from numerous similar disks in a storage cluster.

REFEREED PUBLICATIONS **Practical Design Considerations for Wide Locally Recoverable Codes (LRCs)**

Kadecodi Saurabh, Silas Shashwat, Clausen David, and Merchant Arif. ***FAST 2023 (to appear)***

Tiger: disk-adaptive redundancy without placement restrictions

Kadecodi Saurabh, Maturana Francisco, Athlur Sanjith, Merchant Arif, Rashmi K. V., and Ganger Gregory R. ***OSDI 2022***

WineFS: a hugepage-aware file system for persistent memory that ages gracefully

Kadecodi Rohan, Kadecodi Saurabh, Ponnappalli Soujanya, Shirwadkar Harshad, Ganger Gregory R., Kolli Asheesh, Chidambaram Vijay ***SOSP 2021***

Pacemaker: Avoiding HeART attacks in storage clusters with disk-adaptive redundancy

Kadecodi Saurabh, Maturana Francisco, Suhas Jayaram Subramanya, Yang Junchen, Rashmi K. V. and Ganger Gregory R. ***Highlights presentation in ACM Syster 2022. (arXiv) OSDI 2020***

Cluster storage systems gotta have HeART: improving storage efficiency by exploiting disk-reliability heterogeneity

Kadecodi Saurabh, Rashmi K. V. and Ganger Gregory R. ***(press) FAST 2019***

	Geriatrics: Aging what you see and what you don't see <u>Kadekodi Saurabh</u> , Nagarajan Vaishnavh, Ganger Gregory R. and Gibson Garth A. <i><u>ATC 2018</u></i>
	A case for packing and indexing in cloud file systems <u>Kadekodi Saurabh</u> , Fan Bin, Madan Adit, Gibson Garth A. and Ganger Gregory R. <i><u>HotCloud 2018</u></i>
	Software-Defined Storage for Fast Trajectory Queries using a DeltaFS Indexed Massive Directory Zheng Qing, Amvrosiadis George, <u>Kadekodi Saurabh</u> , Kuchnik Michael, Cranor Chuck, Gibson Garth A., Settlemyer Brad, Grider Gary, and Guo Fan <i><u>PDSW 2017</u></i>
	Caveat-Scriptor: Write Anywhere Shingled Disks <u>Kadekodi Saurabh</u> , Pimpale Swapnil and Gibson Garth A. <i><u>HotStorage 2015</u></i>
	Space Maps in Ext4 <u>Kadekodi Saurabh</u> and Jain Shweta. <u>To be integrated in the mainline Linux kernel.</u> <i><u>OLS 2010</u></i>
UNDER SUBMISSION	<i>Quilt: stitching downsampled IO traces to form full-resolution IO traces</i> <u>Kadekodi Saurabh</u> , Other Authors.
AWARDS AND GRANTS	Helped write a successful NSF grant proposal that received federal funding of \$986,805 (#1956271) on understanding and addressing device-reliability heterogeneity in large-scale distributed storage systems. Nominated by CMU for ACM SIGOPS Best Dissertation Award 2021 for Ph.D. thesis on Disk-Adaptive Redundancy. Best Alumni Research Award 2010 PICT for Space Maps in Ext4. USENIX travel grant - FAST 2015, FAST 2016, FAST 2018, FAST 2019. NSF travel grant - ACM SOSP 2020.
TECHNICAL REPORTS	Compression in Checkpointing and Fault-Tolerant Systems <u>Kadekodi Saurabh</u> <i><u>Technical Report</u></i> VM Co-Migration <u>Kadekodi Saurabh</u> , Shi Chao, Zhang Qingyuan <i><u>Technical Report</u></i> NIC of Time: smart NIC offloading Amos Brandon, <u>Kadekodi Saurabh</u> , Kiyak Serhat <i><u>Technical Report</u></i> Active Databases Avci Besim, <u>Kadekodi Saurabh</u> , Paul Arindam <i><u>Technical Report</u></i> DNA Compression <u>Kadekodi Saurabh</u> <i><u>Technical Report</u></i>
CURRENT WORK	<u>Google</u> - Research Scientist <i>Jun 2022 - present</i> Designing analytical/simulation frameworks, exploring machine learning methods, and drawing insights from massive data sets across various storage systems at Google. Provide critical inputs to inform both short-term tactical actions and long-term technology choices and strategies to help better understand Google's workload and optimize storage design in an increasingly heterogeneous fleet.
RESEARCH EXPERIENCE	<u>Google</u> - Visiting Faculty Researcher <i>Jan 2021 - Jun 2022</i> Generalized disk-adaptive redundancy by removing placement constraints (Tiger). Designing efficient redundancy schemes for cold data. Performed workload analysis of large-scale cluster storage systems. Performed analysis of factors correlating with disk failures in cluster storage systems. <u>Google</u> - Student Researcher <i>Summer 2019 - Dec 2020</i> Analyzed disk failure and reliability data for Google's cluster storage systems. <u>NetApp</u> - Research Intern <i>Spring 2019</i> Analyzed disk failure and reliability data for NetApp's enterprise storage systems spanning over 4.3 million HDDs over 60 makes/models.

Alluxio - PhD Intern*Summer 2017*

Designed and developed a packing and indexing layer in the writeback cache of Alluxio (a distributed user-level file system) to transform arbitrary user workloads to write patterns that are better suited for the cloud. My optimization resulted in $>60\times$ increase in throughput and $>25000\times$ reduction in cost required to ingest data into Amazon S3.

Google - PhD Intern*Summer 2016*

Worked at the intersection of the Storage Analytics and Colossus (distributed file system) teams on analyzing access patterns from Google services and designing a load-balancing strategy to avoid disk hot-spots in data centers with over a hundred thousand disks with heterogeneous capacities and speeds.

Apple - PhD Intern*Summer 2015*

Spent a summer as a part of the file systems team in the CoreOS division.

University of Wisconsin - Madison - Graduate Research Intern*Summer 2014*

Advisor: Prof. Remzi Arpaci-Dusseau

SSD over Infiniband

This study compared the performance between a locally connected SSD and a remotely connected SSD (over Infiniband). Using the lightweight SCSI RDMA protocol (SRP) for communication, we analyzed the costs of accessing remote SSDs and gained insights into enhancing software architectures of next-gen data centers from the storage perspective.

Price of Ext4

Measured the software overhead of the Ext4 file system with the advent of storage devices with microsecond latencies. We threw light on simple optimizations that can be done to remove the bottlenecks observed in sub-modules of Ext4.

WORK

Soft Corner - Software Developer*Oct 2010 - Sep 2012*

EXPERIENCE

PatientScribe

Worked on data reliability, fault tolerance, and audit trails at the clinical level. Developed global data reliability using data analytics. Also developed a distributed data-storage algorithm for seamless remote backup and restore.

Project-X

Designed and developed a custom stackable context-based framework for the social network. Developed a suggester module based on a graph-based traversal of the user connection web to help generate meaningful suggestions based on connected entities. Extended WWW SQL Designer by adding collaboration, versioning, and a chat module to build a collaborative real-time block diagram and discussion tool.

Spring Computing Pvt. Ltd. - Software Engineer*Jun 2009 - Oct 2010***Emdebian Research**

Created a 9 MB JFFS2 filesystem image using Emdebian Crush (an online repository that helps create BusyBox-based root filesystem images) and ported the resulting filesystem image on an ARM board.

Kernel Porting on Embedded Devices

Participated in the full cycle from porting Linux kernel 2.6.30 onto 3 customized ARM boards, testing and tweaking peripherals to adapt to the new kernel and making public releases of the boards.

TEACHING

Lead Teaching Assistant for Storage Systems (CMU 15-746)*Fall 2015***Teaching Assistant** for Distributed Systems (CMU 15-640)*Fall 2016*

Guest lecture at CMU on *Packing and Indexing in Cloud File Systems* in the Advanced Cloud Computing course

Spring 2018

Guest lecture at CMU on *GFS* in the Distributed Systems course

Fall 2016

Guest lecture at CMU on *Caveat Scriptor* in the Storage Systems course

Fall 2015

OUTREACH

Mentored an undergraduate student through the NSF REU program for underrepresented communities.

Summer 2022

Guest lecture at Rutgers University on *Memory & Computers*: making computational memory architectures accessible to a non-specialist audience in the *Pasts and Presents: Recollection, Archives*,

and *Curation* seminar course.

Fall 2018, Spring 2020, Fall 2021

Talk at Shelter Associates on computer science and storage systems research for healthcare and sanitation field workers and administrative staff. Dec 2021

LEADERSHIP **Mentored 16 undergraduate and 16 Masters students' capstone projects** - across 4 areas: building social networks, shingled disk layouts, optimizing for small files in cloud file systems, and incorporating reliability heterogeneity awareness in HDFS. Fall 2010 - present

Co-Hosted Ph.D. research intern - at Google. Built the Quilt system to synthesize realistic disk IO traces by stitching sampled IO traces. Summer 2022

Mentored 2 junior Ph.D. students - who were in their first year while I was a senior Ph.D. student, and continued mentorship after my graduation. Spring 2020 - present

Product Owner / Team Lead of Project-X - led a team of engineers, who in conjunction with a team of undergraduate students built a social network from scratch. Project-X ran successfully in India for about 5 years with a few thousand users and provided a platform for faculty, students, and industry to connect via a social network of undergraduate capstone projects. 2010 - 2012

President of the PICT Art Circle - led the undergraduate drama society and expanded it from 30 to 50 members over two years. 2007 - 2009

Mentored Teaching Assistants - for two iterations of the storage systems course. 2016 - 2017

President of SPICMACAY-CMU - led the *Society for Promoting Indian Classical Music And Culture Among Youth* (CMU student branch) and organized two concerts of internationally acclaimed Indian classical music artists. May 2019 - May 2020

Event Chair (Drama) and elected board member of CMU's Indian Graduate Student Association (IGSA) - established a 30-person team of theater enthusiasts at CMU and performed two one-act plays translated from Marathi to English. 2016 - 2018

INVITED TALKS **Talk at joint USENIX OSDI / ATC** providing an expert overview of selected conference papers in storage systems Jul 2021

Talk at CMU PDL in the summer invited talk series Jul 2021

Talks at Google on *Pacemaker* in the D, Colossus and Storage Analytics teams Aug 2020

Talks at Google on *HeART* in D, Colossus and Storage Analytics teams on *HeART* Jan 2019

Talk at Microsoft Research (India) on *HeART* and disk-adaptive redundancy Dec 2018

Talk at Google on *Geriatrix* in D, Colossus teams Jul 2018

CONFERENCE / WORKSHOP PRESENTATIONS **Presented Tiger at OSDI 2022** Jul 2022

Highlights presentation of Pacemaker at ACM Systor 2022 originally presented at USENIX OSDI 2020. Jun 2022

Presented Pacemaker at OSDI 2020 Nov 2020

Presented HeART at the 1st Toronto Systems Workshop at University of Toronto Oct 2019

Presented HeART at FAST 2019 Feb 2019

Presented Geriatrix at ATC 2018 Jul 2022

Presented Packing and Indexing in Alluxio at HotCloud 2018 Jul 2018

Presented Caveat-Scriptor at HotStorage 2015 Jul 2015

Presented Space Maps at OLS 2010 Jul 2010

SERVICE	External reviewer for ACM Transaction on Storage.	<i>ACM TOS 2022</i>
	External reviewer for ACM Transaction on Computer Systems.	<i>ACM TOCS 2021</i>
	External reviewer for IEEE Transaction on Computers.	<i>IEEE TC 2021</i>
	Area reviewer for the accepted papers in storage systems	<i>OSDI / ATC 2021</i>
	Part of the editorial board for the storage systems area of <u>JSys</u> , the first diamond open-access journal covering all areas of computer systems research.	<i>JSys 2021</i>
	Student reviewer as a part of the shadow program committee	<i>EuroSys 2020</i>
OPEN SOURCED CODE	Geriatric – an aging framework for file systems.	<u>Github</u>
	WineFS – a huge-page friendly, age-resistant persistent memory file system.	<u>Github</u>
	NIC of time – a smart NIC offloading framework.	<u>Github</u>
	Peloton GC – garbage collection for an in-memory database.	<u>Github</u>
STUDENTS MENTORED	Undergraduate Students – Harshad Shirwadkar, Aditya Dani,, Piyush Nimbalkar, Shardul Mangade, Abhijit Jagdale, Ashish Bhumkar, Priti Biyani, Narita Pandhe, Rushikesh Bhome, Ranjithkumar Burgute, Amey Gaidhani, Swapnil Patil, Pavan Ghuge, Sagar Gavitt, Saurabh Wankhade, Austin Ramos	
	Masters Students – Venkata Ravi Chandra Bandlamudi, Swapnil Pimple, Fiona Jane Britto, Omkar-Gawde, Haixin Liu, Xingyu Jin, Jiaan Dai, Jiaqi Zuo, Jiongtao Ye, Sai Kiriti Badam, Xuren Zhou, Arvind Sai Krishnan, Vilas Bhat, Shaobo Guan, Timothy Taeho Kim	
	Ph.D. Students – Suhas Jayaram Subramanya, Sanjith Athlur, Soroush Ghodrati	
REFERENCES	<u>Prof. Gregory Ganger</u> (ganger@ece.cmu.edu)	<i>Carnegie Mellon University</i>
	Jatras Professor of Electrical and Computer Engineering, Professor of Computer Science, Director of the Parallel Data Laboratory	
	<u>Prof. Rashmi K. Vinayak</u> (rvinayak@cs.cmu.edu)	<i>Carnegie Mellon University</i>
	Assistant Professor, Computer Science	
	<u>Prof. Garth Gibson</u> (garth@vectorinstitute.ai)	<i>Vector Institute, Carnegie Mellon University</i>
President and CEO of Vector Institute, Professor of Computer Science and Electrical and Computer Engineering		
<u>Dr. Arif Merchant</u> (aamerchant@google.com)	<i>Google</i>	
Research Scientist and Leader of the Storage Analytics Group		
<u>Prof. Remzi Arpaci-Dusseau</u> (remzi@cs.wisc.edu)	<i>University of Wisconsin-Madison</i>	
Grace Wahba Professor and Chair of Computer Sciences		